

Path-complete positivity of switching systems

F. Forni^{*}, R. M. Jungers^{**}, R. Sepulchre^{***}

^{*} *University of Cambridge, United Kingdom
(f.forni@eng.cam.ac.uk).*

^{**} *Université catholique de Louvain, Belgium
(raphael.jungers@uclouvain.be)*

^{***} *University of Cambridge, United Kingdom
(r.sepulchre@eng.cam.ac.uk)*

Abstract: The notion of path-complete positivity is introduced as a way to generalize the property of positivity from one LTI system to a family of switched LTI systems whose switching rule is constrained by a finite automaton. The generalization builds upon the analogy between stability and positivity, the former referring to the contraction of a norm, the latter referring to the contraction of a cone (or, equivalently, a projective norm). We motivate and investigate the potential of path-positivity and we propose an algorithm for the automatic verification of positivity.

Keywords: Positivity, Path-complete Lyapunov functions, Switching systems, Monotonicity, Perron-Frobenius theory.

1. INTRODUCTION

Positivity is a classical concept of linear system theory. It originates in the many examples of system dynamics whose state variables remain positive along trajectories, and finds its theoretical foundations in Perron-Frobenius theory. In a nutshell, under mild assumptions, the solutions of a positive system converge to a dominant eigendirection in the positive orthant Luenberger (1979). Positivity has known a renewed interest in the recent years for its advantageous computational scalability over general linear systems Rantzer (2015). As a geometric concept, positivity is primarily about the contraction of a cone under the action of a linear map. The positive orthant is a cone of special interest, but Perron-Frobenius theory owes fundamentally to the geometric contraction of a cone more than to an algebraic property of matrices with positive elements.

It is the same contraction property that makes positivity the infinitesimal (or differential) characterization of monotonicity : the order preserving property of a monotone map is equivalent to a positivity property for the linearized map. This geometric viewpoint on positivity is at the root of the differential positivity theory recently introduced in Forni and Sepulchre (2016) to characterize and study the asymptotic properties of nonlinear systems whose trajectories infinitesimally contract a smooth cone field. It has proven quite insightful to think of differential positivity as an analog of differential stability, or contraction analysis. In one case, one studies the contraction of a smooth *norm* field, e.g. a Riemannian metric, while in

the latter case, one studies the infinitesimal contraction of a *cone* field. This insight points to a basic but profound similitude between stability and positivity : two contraction properties, that only differ by the geometric nature of the object that is contracted.

The present paper draws upon this analogy to generalize the concept of positivity from a single matrix (or linear operator) to a family of matrices. Such a generalization has received considerable attention in the context of stability, but much less in the context of positivity. In particular, we focus in the present paper on the recent framework of path-complete Lyapunov analysis, which is a unifying approach to study the stability of a switched system whose switching rule is constrained by a finite automaton. Our goal here is to mimic this framework when the *norm contraction* underlying stability is replaced by a *cone contraction* underlying positivity.

The paper is organized as follows. In Section 2, we briefly recall the notion of positivity, and its links with stability. Then, in Section 3, we naturally draw on this parallel to introduce our main concept: path-complete positivity. In Section 4 we explain what this concept implies in terms of dynamical systems and control, and finally Section 5 touches upon the algorithmic problem of recognizing this property for a given set of matrices.

2. POSITIVITY VERSUS STABILITY

Both stability and positivity are classical notions in linear systems analysis. We review basic notations and terminology and stress the analogy between these two properties in the elementary context of a linear time-invariant (LTI) system $x^+ = Ax$.

^{*} R.J. is currently on sabbatical at UCLA, Department of Electrical Engineering, Los Angeles, USA. His work is supported by the French Community of Belgium and by the IAP network DYSCO. He is a Fulbright Fellow and a FNRS Fellow.

Stability refers to the invariance of a norm, i.e. a *ball* in the state-space. The restriction to *quadratic* norms $|x|_P := \sqrt{x^T P x}$ (where P is a positive definite matrix) is no loss of generality for LTI systems, in which case the invariance condition corresponds to the (Lyapunov) inequality

$$A^T P A - \gamma P \preceq 0, \quad 0 \leq \gamma \leq 1.$$

The case $\gamma = 1$ only ensures invariance (i.e. Lyapunov stability) whereas the case $\gamma < 1$ ensures contraction (i.e. exponential stability). In Lyapunov analysis, the norm $V(x) := x^T P x$ is also called a (quadratic) Lyapunov function.

Fundamentally, positivity is the analog property when the ball is replaced by a cone. In this paper, a cone $\mathcal{K} \subseteq \mathbb{R}^n$ always means a convex pointed solid cone. Recall that a set \mathcal{K} is a convex cone if $\alpha x + \beta y \in \mathcal{K}$ for all $x, y \in \mathcal{K}$ and all $\alpha, \beta > 0$. Pointed means $-\mathcal{K} \cap \mathcal{K} = \{0\}$. \mathcal{K} is solid if it contains n linearly independent vectors.

A linear system is *positive* with respect to \mathcal{K} (in short, \mathcal{K} -positive) if

$$A\mathcal{K} \subseteq \mathcal{K}.$$

Positivity only ensures the invariance condition, while *strict positivity* also enforces *contraction*, by requiring that the boundary of the cone is mapped into the interior of the cone

$$A\mathcal{K} \subseteq \text{int}\mathcal{K}.$$

Positivity has a natural metric characterization based on the Hilbert metric $d_{\mathcal{K}}$ associated to the cone \mathcal{K} .

Definition 1. Bushell (1973a) Given a cone $\mathcal{K} \in \mathbb{R}^n$, the corresponding *Hilbert metric* is given by

$$d_{\mathcal{K}}(x, y) := \log \left(\frac{M_{\mathcal{K}}(x|y)}{m_{\mathcal{K}}(x|y)} \right) \quad \forall x, y \in \mathcal{K}$$

where

$$\begin{aligned} M_{\mathcal{K}}(x|y) &= \inf\{\lambda \mid \lambda y - x \in \mathcal{K}\} = \inf\{\lambda \mid \lambda y \in x + \text{bdr}\mathcal{K}\}; \\ m_{\mathcal{K}}(x|y) &= \sup\{\mu \mid x - \mu y \in \mathcal{K}\} = \sup\{\mu \mid \mu y \in x - \text{bdr}\mathcal{K}\}. \end{aligned}$$

We take $M_{\mathcal{K}}(x|y) = \infty$ if $\forall \lambda > 0, \lambda y \notin x + \mathcal{K}$.

The Hilbert metric is in fact a *distance among rays of the cone*, satisfying the property $d_{\mathcal{K}}(\alpha x, \beta y) = d_{\mathcal{K}}(x, y)$ for any positive scaling α and β . It is therefore a distance in the projective space. In short, contraction of a ball is measured by a norm distance, whereas contraction of a cone is measured by a projective distance. The Hilbert metric characterizes the contraction of a cone in the same way as a Lyapunov function characterizes the contraction of a ball, as shown in the following theorem.

Theorem 1. Bushell (1973a) Consider a matrix $A \in \mathbb{R}^{n \times n}$. If A is \mathcal{K} -positive, then there exists $\gamma < 1$ such that for any $x, y \in \mathcal{K}$

$$d_{\mathcal{K}}(Ax, Ay) \leq \gamma d_{\mathcal{K}}(x, y). \quad (1)$$

Moreover, the smallest γ satisfying the equation above satisfies

$$\gamma = \tanh \frac{1}{4} D_{\mathcal{K}}^A$$

where $D_{\mathcal{K}}^A := \sup_{x, y \in \text{int}\mathcal{K}} d_{\mathcal{K}}(Ax, Ay)$.

Clearly $\gamma < 1$ whenever $D_{\mathcal{K}}^A < \infty$, that is, whenever $A\mathcal{K} \subseteq \text{int}\mathcal{K}$. In what follows we will say that \mathcal{K} is a γ -*contracting* cone for the linear map A whenever (1) holds.

Proving the contraction of a map is a fundamental way of characterizing the existence of a fixed point. Contraction of a ball implies that the iterated map eventually shrinks to a point. This is the essence of Lyapunov theory. Likewise, contraction of a cone implies that the iterated map eventually shrinks to a ray (a point in the projective space). This is the essence of Perron-Frobenius theory.

For a LTI system, both stability and positivity have a spectral characterization. Exponential stability (or contraction) means that all the eigenvalues have a strictly negative real part, while strict positivity (or projective contraction) means that the matrix A has a dominant eigenvector in the interior of the cone.

3. CONTRACTION AND PATH-CONTRACTION

There exists an extensive literature devoted to generalizing the stability of a single matrix (in the sense recalled in the previous section) to a finite (or even compact) set of matrices $A_{\sigma} \in \mathbb{R}^{n \times n}$, $\sigma \in \Sigma := \{1, \dots, N\} \subset \mathbb{N}$. See for instance Liberzon (2003); Jungers (2009) and references therein. One obvious application is the stability analysis of switched systems $x^+ = A_{\sigma} x$ where the update rule is allowed to switch among the considered set of matrices. Drawing upon the analogy stressed above, this section generalizes positivity to a set of matrices.

3.1 Uniform positivity

A straightforward extension with respect to the previous section is to study uniform positivity of a family of matrices with respect to a common cone.

Not surprisingly, strict positivity of each matrix (possibly with respect to different cones) is necessary but not sufficient for uniform strict positivity. And proving the existence of a common invariant cone for a set of linear dynamics is hard. Actually, the existence question is algorithmically undecidable Protasov (2010); Protasov and Voynov (2012); Rodman et al. (2010), very much for the same reasons as its companion question of uniform norm contraction (see Blondel and Tsitsiklis (2000), or (Jungers, 2009, Section 2.2.3)).

It would certainly be of interest to revisit the large body of literature on uniform stability in the light of the analog question of uniform positivity. Even the question of defining a joint projective radius for a family of positive systems in analogy to the ‘joint spectral radius’ defined for a family of stable systems seems valuable and not entirely straightforward. We do not pursue this question in the present paper and leave it for future research.

3.2 Constrained switching systems

Uniform positivity or uniform stability is too conservative of a property for the many applications where the switching rule is not arbitrary. This has long been acknowledged in the literature of switched systems, see for instance Essick et al. (2015); Bliman and Ferrari-Trecate (2003); Lin and Antsaklis (2009), where the permissible sequence of switches is typically modeled by a finite automaton.

Consider a class of switching linear systems represented by

$$x(k+1) = A_{\sigma(k)}x(k) \quad (2)$$

where $\sigma \in \Sigma := \{1, \dots, N\} \subset \mathbb{N}$ and each A_σ is a $n \times n$ matrix. For a switching signal $\sigma(\cdot) : \mathbb{N} \rightarrow \Sigma$ and any initial condition $x_0 \in \mathbb{R}^n$, the unique solution $x(\cdot) : \mathbb{N} \rightarrow \mathbb{R}^n$ of (2) is called a *trajectory* of the system. We say that the system is a *constrained switching system* if the sequences $\sigma(0)\sigma(1)\dots$ generated by the switching signal $\sigma(\cdot)$ belongs to a regular language L_r .

Thus, $\sigma(\cdot)$ is generated by any finite-state automaton (Q, Σ, δ) that accepts the same regular language L_r , where Q is the set of states, Σ is the alphabet and $\delta \subseteq Q \times \Sigma \times Q$ is the transition relation. We say that such an automaton is *path-complete* to emphasize the fact that its paths capture a complete description of the allowed behaviors of the switching signal. We will denote any labeled transition by the compact notation $i \xrightarrow{\sigma} j \in \delta$. A finite sequence of transitions from i to j will be represented by $i \xrightarrow{\sigma_1 \dots \sigma_r} j$.

The complexity of the switching behavior is modulated by the automaton. An example is in Figure 1. *Arbitrary switches* between two matrices A_0 and A_1 are easily captured by the automaton on the left. In contrast, the automaton on the right enforces a switching behavior with a strict alternation between 0 and 1. A mixed situation is provided by the automaton in the middle, whose switches sequences allow for any repetition of 1 separated by isolated zeros.




Fig. 1. Automata with different path restrictions.

The case of unconstrained switches is typical of robust analysis where parametric uncertainties are modeled via nondeterministic switches among a family of linear systems Liberzon (2003). Constrained switches arise from literature on hybrid/cyber-physical systems Essick et al. (2015); Bliman and Ferrari-Trecate (2003); Lin and Antsaklis (2009). In constrained switching systems, specific sequences of operations are captured by suitable branches of the automaton. Restrictions on paths could be used to model forms of ergodicity in the sequence of matrix operations, or to model the alternation between periods of local/isolated operations and periods of collective computations.

3.3 Path-complete Lyapunov functions

Since the nineties, several methods have been proposed for the stability analysis of switched systems with or without restrictions on the switching rules Bliman and Ferrari-Trecate (2003); Daafouz et al. (2002); Essick et al. (2014); Branicky (1998). We briefly summarize the recently proposed framework of path-complete Lyapunov functions, Ahmadi et al. (2014), that provides a unifying approach, and generalizes these techniques.

Definition 2. Consider a constrained switching system and let (Q, Σ, δ) be any path-complete automaton. A *path-complete Lyapunov function* is a multiple Lyapunov

function given by a finite set of homogeneous positive definite functions $(V_i)_{i \in Q}$, $V_i : \mathbb{R}^n \mapsto \mathbb{R}^+$, such that

$$V_j(A_\sigma x) \leq \gamma V_i(x).$$

for each transition $i \xrightarrow{\sigma} j \in \delta$ and each $x \in \mathbb{R}^n$.

The reason of this definition lies in the following theorem.

Theorem 2. (Ahmadi et al. (2014)).

Consider a constrained switching system and let (Q, Σ, δ) be any path-complete automaton. The existence of a *path-complete Lyapunov function* for $\gamma = 1$ is a valid criterion for the stability of the switching system. Asymptotic stability requires $0 \leq \gamma < 1$.

We remark that the regular language that constrains the switches of a constrained switching system can be generated by infinitely many automata and each of these provides a different set of path-complete Lyapunov functions. The selection of a suitable automaton is a degree of freedom in path-complete analysis. The number of states of the automaton allows to balance the complexity of the multiple Lyapunov function with the computational efficiency.

3.4 Path-complete positivity

We follow the approach of path-complete Lyapunov functions to define the corresponding notion for positive systems. Once again, the key step is to substitute cones to norms.

Definition 3. Consider a constrained switching system and let (Q, Σ, δ) be any path-complete automaton for this system. The constrained switching system is *path-complete positive* with respect to the set of cones

$$\bar{\mathcal{K}} := \{\mathcal{K}_q \mid q \in Q\}$$

if¹

$$A_\sigma \mathcal{K}_i \subseteq \mathcal{K}_j$$

for each transition $i \xrightarrow{\sigma} j \in \delta$. *Strict path-complete positivity* further requires that

$$A_\sigma \mathcal{K}_i \subseteq \text{int} \mathcal{K}_j$$

for each transition $i \xrightarrow{\sigma} j \in \delta$.

The definition above reduces to positivity when each cone in the set $\bar{\mathcal{K}}$ is identical. Path-complete positivity is a proper generalization of positivity: Example 1 below discusses the case of a path-complete positive switching system that cannot be positive with respect to a common cone.

Example 1. Consider the constrained switching system $x^+ = A_\sigma x$ with

$$A_0 = \begin{bmatrix} 5 & 0 \\ 0 & 1 \end{bmatrix} \quad A_1 = \begin{bmatrix} 1 & 0 \\ 0 & 3 \end{bmatrix};$$

and suppose that the automaton in Figure 2 is path complete.

The system cannot be strictly positive with respect to a common cone since the dominant eigenvector e_1 of the matrix A_0 is a non-dominant eigenvector of the other matrix A_1 and viceversa. It turns out that the system is

¹ Each \mathcal{K}_i is a pointed, convex, solid cone.




Fig. 2. One of the automata generating σ in Example 1.

strictly path-complete positive with respect to the family of cones $\bar{\mathcal{K}} := \{\mathcal{K}_0, \mathcal{K}_1\}$ where

$$\begin{aligned}\mathcal{K}_0 &:= \{x_1 \geq 0, |x_2| \leq x_1\} \\ \mathcal{K}_1 &:= \{x_1 \geq 0, |x_2| \leq x_1/4\}.\end{aligned}$$

One can check that the path-complete inclusions are satisfied with such values of $\mathcal{K}_0, \mathcal{K}_1$: Indeed, following the automaton paths, any $(x_1, x_2) \in \mathcal{K}_0$ is mapped into $(x_1^+, x_2^+) \in \text{int}\mathcal{K}_1$ by A_0 : $A_0\mathcal{K}_0 = \{x_1^+ = 5x_1 \geq 0, |x_2^+| = |x_2| \leq x_1 = x_1^+/5 < x_1^+/4\} \subseteq \text{int}\mathcal{K}_1$. In a similar way, any $(x_1, x_2) \in \mathcal{K}_1$ is mapped into $(x_1^+, x_2^+) \in \text{int}\mathcal{K}_1$ by A_0 : $A_0\mathcal{K}_1 = \{x_1^+ = 5x_1 \geq 0, |x_2^+| = |x_2| \leq x_1/4 = x_1^+/20 < x_1^+/4\} \subseteq \text{int}\mathcal{K}_1$; and any $(x_1, x_2) \in \mathcal{K}_1$ is mapped into $(x_1^+, x_2^+) \in \text{int}\mathcal{K}_0$ by A_1 : $A_1\mathcal{K}_1 = \{x_1^+ = x_1 \geq 0, |x_2^+| = |3x_2| \leq \frac{3}{4}x_1 = \frac{3}{4}x_1^+ < x_1^+\} \subseteq \text{int}\mathcal{K}_0$. \square

The definition of path-complete positivity suggests that one of the advantages of path-complete positivity is that a temporary “excess” of contraction can be “stored” by narrowing cones. A temporary “lack” of contraction can be “balanced” by widening cones. The example shows that this approach can be effective on finite paths: weak contraction at some steps is balanced by the excess of contraction at some other steps.

4. THE ASYMPTOTIC BEHAVIOR OF PATH-POSITIVE SYSTEMS

The connection between positivity and projective contraction of the Hilbert metric is now generalized to path-complete positive systems.

Theorem 3. Consider a constrained switching system, let (Q, Σ, δ) be any path-complete automaton, and suppose that the constrained switching system is *path positive* with respect to the set of cones $\bar{\mathcal{K}} := \{\mathcal{K}_q \mid q \in Q\}$. Then, there exists $0 \leq \gamma \leq 1$ such that, for any transition $i \xrightarrow{\sigma} j$ of the automaton,

$$d_{\mathcal{K}_j}(A_\sigma x, A_\sigma y) \leq \gamma d_{\mathcal{K}_i}(x, y) \quad \forall x, y \in \mathcal{K}_i. \quad (3)$$

Furthermore, *strict path positivity* guarantees $0 \leq \gamma < 1$.

Proof: Following the proof argument for Theorem 3.1 in Bushell (1973a), one shows that path positivity guarantees $m_{\mathcal{K}_i}(x|y) \leq m_{\mathcal{K}_j}(A_\sigma x|A_\sigma y) \leq M_{\mathcal{K}_j}(A_\sigma x|A_\sigma y) \leq M_{\mathcal{K}_i}(x|y)$ for each transition $i \xrightarrow{\sigma} j \in \delta$, which directly implies (3) for $\gamma = 1$.

For strict path positivity (3) with $0 \leq \gamma < 1$ follows by the proof argument of Theorem 3.2 in Bushell (1973a). For instance, for each $q \in Q$, define the oscillation $\text{osc}_{\mathcal{K}_q}(x|y) := M_{\mathcal{K}_q}(x|y) - m_{\mathcal{K}_q}(x|y)$. Theorems 4 and 5 in Bauer (1965) show that $\text{osc}_{\mathcal{K}_j}(A_\sigma x|A_\sigma y) \leq N_{ij}(A_\sigma) \text{osc}_{\mathcal{K}_i}(x|y)$ for each $i \xrightarrow{\sigma} j \in \delta$, where the oscillation ratio $0 \leq N_{ij}(A_\sigma) < 1$ if $A_\sigma \mathcal{K}_i \subseteq \text{int}\mathcal{K}_j$. This result is well known for positive operators from a cone into itself. The proof argument in Bauer (1965) extends to the case of positive operators between two different cones. Finally, using the proof argument of Lemma 3 in Bushell (1973b) one shows that

$$d_{\mathcal{K}_j}(A_\sigma x|A_\sigma y) \leq N_{ij}(A_\sigma) d_{\mathcal{K}_i}(x|y) \text{ for each } i \xrightarrow{\sigma} j \in \delta. \quad \square$$

Thus, $\gamma := \max_{i \xrightarrow{\sigma} j \in \delta} N_{ij}(A_\sigma) < 1$.

At each transition $i \xrightarrow{\sigma} j$ strict positivity guarantees that the linear map A_σ is a contraction on the rays of the cones, in the sense of the adapted Hilbert metrics $d_{\mathcal{K}_i}$. It is easy to prove, by induction, that any pair $(x(\cdot), y(\cdot))$ of trajectories of the system associated to the same switching signal $\sigma(\cdot)$ and such that $x(0), y(0) \in \mathcal{K}_{q(0)} \setminus \{0\}$ satisfy

$$\lim_{k \rightarrow \infty} \left| \frac{x(k)}{|x(k)|} - \frac{y(k)}{|y(k)|} \right| = 0. \quad (4)$$

Equation (4) makes clear that a strictly positive system asymptotically ‘forgets’ its initial condition, as it converges to a unique steady state solution in the projective space, for every switching signal.

Note that the projective contraction property does not enforce convergence to a fixed point. For example, a straightforward consequence of the theorem is that each cyclic path $q \xrightarrow{\sigma_1 \dots \sigma_r} q$ defines a corresponding path-dependent Perron-Frobenius eigenvalue and eigenvector, $\lambda_{\sigma_1 \dots \sigma_k}$ and $v_{\sigma_1 \dots \sigma_r}$, such that

$$A_{\sigma_r} \dots A_{\sigma_1} v_{\sigma_1 \dots \sigma_r} = \lambda_{\sigma_1 \dots \sigma_k} v_{\sigma_1 \dots \sigma_r}$$

(since $\bar{A} := A_{\sigma_1}, \dots, A_{\sigma_r}$ is necessarily a strictly positive matrix). Denoting rays by $[x] := \{\lambda x \mid \lambda > 0\}$, a simple permutation of indices shows that $[v_{\sigma_2 \dots \sigma_r \sigma_1}] = [A_{\sigma_1} v_{\sigma_1 \dots \sigma_r}]$, $[v_{\sigma_3 \dots \sigma_r \sigma_1 \sigma_2}] = [A_{\sigma_2} v_{\sigma_2 \dots \sigma_r \sigma_1}]$, and so on. Indeed, all the path-dependent Perron-Frobenius eigenvectors on a cyclic path define an invariant sequence of rays. Such sequence is also an attractor of the system. Thus, trajectories along these cycles either converge to zero or to a limit cycle of r rays. In that sense, path-positivity retains the fundamental contraction property of a positive system.

5. ALGORITHMS FOR DECIDING POSITIVITY

Testing the existence of a common invariant or contractive cone is hard Protasov (2010). In fact Protasov proved that the question of whether a set of matrices has an invariant cone is Turing-undecidable. His construction suggests that the question is hard when the matrices share a common invariant linear subspace. For matrices that do not share a common invariant subspace, we algorithmically test whether a given set of matrices has a common γ -contracting cone, for a given contraction ratio $0 < \gamma < 1$. We only discuss the algorithm in the case of uniform positivity and leave for future work a generalization to path-complete positivity.

5.1 Basic test

A single matrix admits a contracting cone if and only if it has a leading eigenvector. An obvious necessary condition for uniform strict positivity w.r.t. a common cone \mathcal{K} is therefore that each system A_σ has a leading eigenvector. We introduce a corresponding splitting of the state-space, which relies on the eigenstructure of A_σ .

Definition 4. For any positive matrix A_σ , we define the *invariant splitting* of \mathbb{R}^n ($\mathcal{V}_\sigma, \mathcal{N}_\sigma$) as the pair of two

A_σ -invariant subspaces of dimension 1 and $n-1$ respectively. \mathcal{V}_σ is defined as the span of the Perron Frobenius eigenvector of A_σ . \mathcal{N}_σ is the unique $n-1$ invariant subspace for A_σ such that $\mathcal{V}_\sigma \cap \mathcal{N}_\sigma = \{0\}$ (for example \mathcal{N}_σ could be defined by the columns of the coordinate transformation that brings A_σ into its real Jordan form).

An elementary necessary condition is as follows.

Proposition 4. If a cone \mathcal{K} is invariant for the matrix A_σ , then necessarily

$$\mathcal{K} \cap \mathcal{N}_\sigma = \{0\}.$$

Corollary 5. (Basic test). If a set of matrices $\mathcal{M} = \{A_\sigma\}$ share a common contracting cone, then they all have a strictly dominant eigenvalue, and the corresponding eigenvector v_σ does not belong to any $\mathcal{N}_{\sigma'}$ for any $\sigma' \neq \sigma$.

5.2 Inner bound

The basic idea of our algorithm below is to start from an inner bound, and proceed by forward iteration (i.e. apply our matrices to this inner bound) in order to enlarge it. For the initial inner bound, one can start with the convex hull of the leading eigenvectors of the matrices, which must be in any invariant cone. Given the set of leading eigenvectors $\{v_i\}$, it is not clear however whether to use v_i or $-v_i$ in the initial inner bound. We resolve this choice as follows: pick any matrix A_σ and define w as the normal vector to the invariant subspace \mathcal{N}_σ . Then for each leading eigenvector of the matrices A_i , pick the orientation v_i such that $v_i^T w > 0$. We formalize the argument in the following proposition.

Proposition 6. (The orientation trick). Suppose that $A_1, A_2 \in \mathbb{R}^{n \times n}$ have a common contracting cone \mathcal{K} , and note v_1, v_2 the leading eigenvectors of A_1, A_2 . Suppose without loss of generality that $v_1 \in \mathcal{K}$. Then, with the notations of Definition 4, v_2 is also in \mathcal{K} if and only if $(w_1^T v_1) \cdot (w_1^T v_2) > 0$, where w_1 is the normal vector to \mathcal{N}_1 .

Proof: If: one has either $v_2 \in \mathcal{K}$, or $-v_2 \in \mathcal{K}$. Now, if $(w_1^T v_1) \cdot (w_1^T v_2) > 0$, it means that $(w_1^T v_1) \cdot (w_1^T (-v_2)) < 0$, and then there exist $\alpha, \beta > 0$ such that $\alpha v_1 + \beta (-v_2) \in \mathcal{N}_1$, and thus $(-v_2)$ cannot belong to \mathcal{K} .

Only if: Suppose by contradiction that $(w_1^T v_1) \cdot (w_1^T v_2) < 0$. Then there exist $\alpha, \beta > 0$ such that $\alpha v_1 + \beta v_2 \in \mathcal{N}_1$, and this contradicts v_1, v_2 being in \mathcal{K} , because $\mathcal{K} \cap \mathcal{N}_1 = \{0\}$. \square

By construction, the convex hull of leading eigenvectors (selected with the proper orientation) is an invariant cone. It thus provides an inner bound for the contracting cone \mathcal{K} . However, the following proposition shows that this cone cannot be contracting, even if there exists a contracting cone.

Proposition 7. Consider a set of matrices $\mathcal{M} \in \mathbb{R}^{n \times n}$, and the set of leading eigenvectors $v_i \in \mathbb{R}^n$ of the matrices in \mathcal{M} . Suppose that

$$\mathcal{K}_0 = \text{conic-hull} \bigcup_{A \in \mathcal{M}^*, v_i} A v_i$$

is a closed convex pointed cone². Then, \mathcal{K}_0 is an invariant cone, but not a γ -contracting cone for any $\gamma < 1$.

² \mathcal{M}^* is the set of all the products of matrices of \mathcal{M} .

Proof: It is obvious that \mathcal{K}_0 is invariant by definition of \mathcal{K}_0 . Now, let us suppose by contradiction that $\mathcal{M}\mathcal{K}_0 \subset \text{int } \mathcal{K}_0$. Since every v_i is an eigenvector of some matrix $A \in \mathcal{M}$, conic-hull $\{v_i\}$ is not contracting. Thus, there exists some $x^* \in \mathcal{K}_0 \setminus \text{int } \mathcal{K}_0$, $x^* \notin \text{conic-hull } \{v_i\}$. By definition of \mathcal{K}_0 , for any $\epsilon > 0$, there is a $x' \in \mathcal{K}_0$, $A \in \mathcal{M}$ such that $|Ax' - x^*| < \epsilon$. This is in contradiction with the fact that $\mathcal{M}\mathcal{K}_0 \subset \text{int } \mathcal{K}_0$, $\mathcal{M}\mathcal{K}_0$ being a finite union of closed sets strictly contained in \mathcal{K}_0 . \square

Example 2. Consider the set of matrices $\mathcal{M} = \{A_1, A_2\}$.

$$A_1 = \begin{bmatrix} 2 & 0 \\ 1.65 & 0.5 \end{bmatrix} \quad A_2 = \begin{bmatrix} 2 & 0 \\ 1.3636 & 0.5 \end{bmatrix}.$$

The leading eigenvectors are $v_1 = [1, 1]^T$ and $v_2 = [1, 1]^T$. \mathcal{M} possesses an invariant cone, which is

$$\text{conic-hull } \{v_i\}.$$

However, the cone $\mathcal{K}_\epsilon = \text{conic-hull } \{[\epsilon, 1]^T, [1, 0]^T\}$, for $\epsilon > 0$ small enough, is contracting (take for instance $\epsilon = 0.1$). If an algorithm proceeds by forward iteration starting from conic-hull $\{v_i\}$, it remains stuck in the cone delimited by these two vectors:

$$\mathcal{K}_0 = \text{conic-hull} \bigcup_{A \in \mathcal{M}^*, v_i} A v_i = \text{conic-hull } \{v_i\}.$$

That is, \mathcal{K}_0 is an invariant cone, but not contracting. It is however included in the contracting cone \mathcal{K}_ϵ . \square

5.3 Algorithm with guaranteed termination and accuracy

In any contracting cone, Theorem 1 implies a uniform upper bound $D_{\mathcal{K}}^A$ on the distance between two points in $A\mathcal{K}$. This bound on the distance is useful to build a *larger inner bound on the contracting cone \mathcal{K}* : indeed, an upper bound on the distance between two points translates geometrically into a lower bound on the distance between any of these points and the boundary of the considered cone (see Definition 1 of the Hilbert metric, and the proof of Lemma 8 below). Thus, we can leverage this information in order to inflate the cone, by ‘pushing the boundaries’ of our inner bound. We formalize this in the next lemma:

Lemma 8. Let \mathcal{K} be a γ -contracting cone for a set of linear maps \mathcal{M} , and take a matrix $A \in \mathcal{M}$. For any $x, y \in \mathcal{K}$, consider $x' = Ax$, $y' = Ay$. Suppose that $x' - y' \notin \mathcal{K}$; then, for any $\rho \geq \exp(D_{\mathcal{K}}^A)$ we have that

$$y' + \frac{1}{\rho - 1}(y' - x') \in \mathcal{K}.$$

Proof: From Definition 1 we have $0 \leq m(Ax|Ay) < 1$ since $Ax - Ay \notin \mathcal{K}$. Thus,

$$\begin{aligned} M(Ax|Ay) &= \frac{M(Ax|Ay)}{m(Ax|Ay)} m(Ax|Ay) \\ &\leq \exp(D_{\mathcal{K}}^A) m(Ax|Ay) \\ &\leq \exp(D_{\mathcal{K}}^A). \end{aligned}$$

Furthermore, $M(Ax|Ay)Ay - Ax \in \mathcal{K}$, thus $\rho Ay - Ax \in \mathcal{K}$ since $\rho \geq M(Ax|Ay)$. Finally, writing $\rho Ay - Ax + Ay - Ay \in \mathcal{K}$, we obtain $Ay + \frac{1}{\rho - 1}(Ay - Ax) \in \mathcal{K}$. \square

Lemma 8 provides a way to widen any inner bound of \mathcal{K} in such a way that the widened cone is still a subset of \mathcal{K} .

Indeed, if an inner bound is not contractive, we can use the lemma to widen its boundary slightly outwards before pursuing the forward iteration algorithm. Lemma 8 and Lemma 9 below are at the core of Algorithm 1, which decides in finite time whether a given set of matrices has a common γ -contracting cone, as clarified in Theorem 10.

Lemma 9. Let \mathcal{K} be a cone in \mathbb{R}^n , $x, y \in \text{int } \mathcal{K}$, and suppose that $d_{\mathcal{K}}(x, y) > 0$. Then, for any $(n-1)$ -dimensional hyperplane H such that $H \cap \mathcal{K} = \{0\}$, there exists a $\lambda > 0$ such that $y - \lambda x \in H$.

Proof: For λ very small, we have $y - \lambda x \in \mathcal{K}$; for λ very large, we have $y - \lambda x \in -\mathcal{K}$. Thus, by continuity, there must be a λ such that $y - \lambda x \in H$. \square

In the next theorem, we suppose that the matrices do not have zero eigenvalues, nor common invariant subspace. These are technical assumptions that hold for generic matrices.

Theorem 10. Consider a set of positive matrices \mathcal{M} with nonzero eigenvalues and no common invariant subspace. Given a contraction ratio γ , Algorithm 1 decides in finite time whether the set of matrices has a common γ -contracting cone.

- It returns a γ -contracting cone provided that such a cone exists.
- If there is no γ -contracting cone, it returns ‘NO’, or a δ -invariant cone, for $\gamma < \delta < 1$, if it has found one.

Data: A set of matrices \mathcal{M} , a number $\gamma \in (0, 1)$

Result: Outputs YES if the set of matrices has a γ -contracting cone (and returns a description of the invariant cone). If the set of matrices has a strictly invariant cone, but no γ -contracting cone, it may return NO, or a δ -invariant cone for some $\delta > \gamma$.

begin

$\mathcal{K}_0 = \text{conic-hull}\{v \mid v : \text{leading eigenvector of } A \in \mathcal{M}\}$

% vectors v are picked according to Proposition 6

$\rho \geq \exp(4 \tanh^{-1}(\gamma))$

$t = 0$

while true do

if $\mathcal{K}_t \cap \mathcal{N}_\sigma \neq \{0\}$ *for some* \mathcal{N}_σ **then**

 Output NO; Exit.

end

if \mathcal{K}_t *is strictly invariant* **then**

 Output YES; Exit.

end

$\mathcal{K}_{t+1} = \text{conic-hull}\{\bigcup_{A \in \mathcal{M}} A\mathcal{K}_t \cup \mathcal{K}_t\}$

 For all $A \in \mathcal{M}$, for all vertices $y, x \in \mathcal{K}_t$

 such that $(Ay - Ax) \in \mathcal{N}_\sigma$ for some $\sigma \in \mathcal{M}$,

$\mathcal{K}_{t+1} = \text{conic-hull}\{\mathcal{K}_{t+1} \cup \{Ay + (Ay - Ax)/(\rho - 1)\}\}$

$t = t + 1$;

end

end

Algorithm 1. An algorithm for deciding joint positivity

Proof: The algorithm iteratively computes inner bounds \mathcal{K}_t for \mathcal{K} . We first prove that indeed \mathcal{K}_t are valid inner

bounds (provided that there indeed exists a contracting cone \mathcal{K}). We then prove that one of these inner bounds \mathcal{K}_t must be contracting for some t (and not only invariant). Thus, the algorithm will terminate with an effective contracting cone.

The algorithm starts with \mathcal{K}_0 as a first inner bound (if \mathcal{K}_0 is not a convex pointed cone, then one directly concludes that the set of matrices does not have a common invariant cone). Thus, suppose that \mathcal{K}_0 is a valid inner bound by Proposition 7. Now, at every step, with an initial inner bound \mathcal{K}_t , the algorithm performs two operations. First, it takes the union of \mathcal{K}_t with all the images of this set

$$\mathcal{K}_{t+1} = \text{conic-hull}\left\{\bigcup_{A \in \mathcal{M}} A\mathcal{K}_t \cup \mathcal{K}_t\right\},$$

which is clearly still an inner bound, by definition of a contracting cone. Then, for any two points $x, y \in \mathcal{K}_t$, it adds the point $\{Ay + (Ay - Ax)/(\rho - 1)\}$ to \mathcal{K}_{t+1} .

For practical efficiency, the true algorithm can only do it for points Ay, Ax that are vertices of the new inner bound. Also, one has to scale x in order to ensure the condition $(Ay - Ax) \in \mathcal{N}_\sigma$, but, provided that Ax and Ay are not parallel, this is always possible by Lemma 9 above. Finally note that there must always be an $Ay \in \mathcal{K}_{t+1} \setminus \text{int } \mathcal{K}_{t+1}$ (if not, \mathcal{K}_{t+1} would be contracting) and Ax non-aligned with Ay (because the matrix A has nonzero eigenvalues). In turn, the condition $(Ay - Ax) \in \mathcal{N}_\sigma$ implies that $Ay - Ax \notin \mathcal{K}$ (Proposition 4), and we can apply Lemma 8. These new added points are guaranteed to be in the invariant cone \mathcal{K} , by Lemma 8, and this proves that \mathcal{K}_{t+1} is still an inner bound.

We now prove that this procedure generates a contracting cone after a finite number of steps (if there exists one). Suppose the contrary. Then, the inner bounds \mathcal{K}_t converge towards a cone \mathcal{K}_∞ which is invariant, but not contracting. Consider a vertex z of \mathcal{K}_∞ , which is such that $Az \in \mathcal{K}_\infty \setminus \text{int } \mathcal{K}_\infty$. That is, Az is in the boundary of \mathcal{K}_∞ . This implies that the inflating step in the algorithm is such that for all $x \in \mathcal{K}_\infty$, Ax is in the same face of \mathcal{K}_∞ as Az (because in the opposite case, the inflation step would ‘push’ Az out of \mathcal{K}_∞ , and Az would not be a vertex anymore). \mathcal{K}_∞ being of nonempty interior (because the matrices have no nontrivial invariant subspace), this implies that A has zero eigenvalues, a contradiction.

In conclusion, the algorithm cannot converge to a non contracting invariant cone. Thus, if there exists a contracting \mathcal{K} , since \mathcal{K}_t are valid inner bounds (i.e. contained in \mathcal{K}), they will either converge to \mathcal{K} , or the algorithm will stop before (having found another contracting cone). If, on the other hand, there is no invariant cone, the ‘inner bounds’ will keep growing until they intersect some \mathcal{N}_σ , and the algorithm will stop, concluding that there is no γ -contracting cone. \square

6. CONCLUSIONS AND FURTHER DIRECTIONS.

In this work, we have introduced the concept of path-complete positivity, which generalizes the notion of positivity. We showed that this notion can be useful, for instance for (constrained) switching systems, for which we provide an example of system which is not positive, but yet, is path-complete positive. We showed that path-complete positive systems inherit much of the nice proper-

ties of positive systems, and we sketched an algorithm to decide whether a switching system has an invariant cone. Our algorithm is inspired from the similar, and much more studied, problem of proving stability for switching systems. It proceeds by *forward propagation*, which is a well-known technique for proving stability of a switching system. However, the positivity problem is more tricky, for several reasons: first, contrary to the stability problem, one cannot take an arbitrary norm for initializing a forward propagation procedure. Second, the forward iteration converges by essence to a *non contracting* invariant cone, forcing us to introduce an ‘inflation procedure’ in order to generate a contracting cone.

We believe that path-complete positivity opens a number of directions worth exploring. The first step is to decide path-complete positivity for a switching system, as mentioned above. In a second time, we plan to generalize this notion to more general dynamical systems, and link it to the notion of differential positivity.

Path-complete positivity may also prove useful for computational goals, even for systems that *do have a common contracting cone*, that is, that are positive. Just like path-complete stability has been used as a proxy to prove stability in the control literature, it might be easier to prove path-complete positivity than to prove positivity.

REFERENCES

- Ahmadi, A.A., Jungers, R.M., Parrilo, P.A., and Roozbehani, M. (2014). Joint spectral radius and path-complete graph lyapunov functions. *SIAM Journal on Control and Optimization*, 52(1), 687–717.
- Bauer, F. (1965). An elementary proof of the hopf inequality for positive operators. *Numerische Mathematik*, 7(4), 331–337. doi:10.1007/BF01436527.
- Bliman, P.A. and Ferrari-Trecate, G. (2003). Stability analysis of discrete-time switched systems through lyapunov functions with nonminimal state. In *Proceedings of IFAC Conference on the Analysis and Design of Hybrid Systems*, 325–330.
- Blondel, V.D. and Tsitsiklis, J.N. (2000). The boundedness of all products of a pair of matrices is undecidable. *Systems & Control Letters*, 41(2), 135–140.
- Branicky, M.S. (1998). Multiple lyapunov functions and other analysis tools for switched and hybrid systems. *IEEE Transactions on automatic control*, 43(4), 475–482.
- Bushell, P. (1973a). Hilbert’s metric and positive contraction mappings in a Banach space. *Archive for Rational Mechanics and Analysis*, 52(4), 330–338. doi:10.1007/BF00247467.
- Bushell, P. (1973b). On the projective contraction ratio for positive linear mappings. *Journal of the London Mathematical Society*, s2-6(2), 256–258. doi:10.1112/jlms/s2-6.2.256.
- Daafouz, J., Riedinger, P., and Iung, C. (2002). Stability analysis and control synthesis for switched systems: a switched lyapunov function approach. *IEEE Transactions on Automatic Control*, 47(11), 1883–1887.
- Essick, R., Lee, J.W., and Dullerud, G.E. (2014). Control of linear switched systems with receding horizon modal information. *IEEE Transactions on Automatic Control*, 59(9), 2340–2352.
- Essick, R., Philippe, M., Dullerud, G., and Jungers, R.M. (2015). The minimum achievable stability radius of switched linear systems with feedback. In *IEEE 54th Annual Conference on Decision and Control (CDC)*, 4240–4245. IEEE.
- Forni, F. and Sepulchre, R. (2016). Differentially positive systems. *IEEE Transactions on Automatic Control*, 61(2), 346–359. doi:10.1109/TAC.2015.2437523.
- Jungers, R. (2009). *The joint spectral radius: Theory and applications*. Lecture Notes in Control and Information Sciences 385. Springer-Verlag Berlin Heidelberg, first edition.
- Liberzon, D. (2003). *Switching in Systems and Control*. Systems & Control. Birkhäuser, first edition.
- Lin, H. and Antsaklis, P.J. (2009). Stability and stabilizability of switched linear systems: a survey of recent results. *IEEE Transactions on Automatic control*, 54(2), 308–322.
- Luenberger, D. (1979). *Introduction to Dynamic Systems: Theory, Models, and Applications*. Wiley, 1 edition.
- Protasov, V. (2010). When do several linear operators share an invariant cone? *Linear Algebra and its Applications*, 433(4), 781 – 789. doi: <http://dx.doi.org/10.1016/j.laa.2010.04.006>.
- Protasov, V. and Voynov, A. (2012). Sets of nonnegative matrices without positive products. *Linear Algebra and its Applications*, 437(3), 749 – 765. doi: <http://dx.doi.org/10.1016/j.laa.2012.02.029>.
- Rantzer, A. (2015). Scalable control of positive systems. *European Journal of Control*, 24, 72 – 80.
- Rodman, L., Seyalioglu, H., and Spitkovsky, I. (2010). On common invariant cones for families of matrices. *Linear Algebra and its Applications*, 432(4), 911 – 926. doi: <http://dx.doi.org/10.1016/j.laa.2009.10.004>.